

# A Game Theoretical Extended Vickrey Auction Mechanism for Replicating Data in Large-scale Distributed Computing Systems

Samee Ullah Khan and Ishfaq Ahmad  
 Department of Computer Science and Engineering  
 University of Texas at Arlington

*Abstract* — This paper addresses the problem of fine-grained data replication on a set of Internet sites using the extended form of Vickrey auction called the  $N+1^{\text{st}}$  price auction. Specifically, we present an adaptive auction mechanism for replication of objects in a distributed system. The mechanism is adaptive in the sense that it changes the replica schema of the objects by continuously moving the schema towards an optimal one, while ensuring object concurrency control. The mechanism is experimentally evaluated against three well-known techniques from the literature: greedy, branch and bound, and genetic algorithms. The experimental results reveal that the proposed approach outperforms the three techniques in both the execution time and solution quality.

**Keywords:** Data replication, auctions, allocation, pricing, resource management.

## 1 Introduction

Fine-grained (object based) replica schemas determine how many replicas of each objects are created, and to which sites they are assigned. These schemas critically affect the performance of the distributed computing system (e.g. the Internet), since reading an object locally is less costly than reading it remotely [13]. Therefore, in a read intensive network an extensive replica schema is required. On the other hand, an update of an object is written to all, and therefore, in a write intensive network a constricted replica schema is required. In essence replica schemas

TABLE I  
 NOTATIONS AND THEIR MEANINGS.

Symbol	Meaning
$M$	Total number of sites in the network.
$N$	Total number of objects to be replicated.
$O_k$	$k$ -th object.
$o_k$	Size of object $k$ .
$S_i$	$i$ -th site.
$s_i$	Size of site $i$ .
$r_k^i$	Number of reads for object $k$ from site $i$ .
$R_k^i$	Aggregate read cost of $r_k^i$ .
$w_k^i$	Number of writes for object $k$ from site $i$ .
$W_k^i$	Aggregate write cost of $w_k^i$ .
$NN_k^i$	Nearest neighbor of site $i$ holding object $k$ .
$c(i,j)$	Communication cost between sites $i$ and $j$ .
$P_k$	Primary site of the $k$ -th object.
$R_k$	Replication schema of object $k$ .
$C_{overall}$	Total overall data transfer cost.

are strongly dependent upon the read and write patterns for each object [3].

Recently, a few approaches on replicating data objects over the Internet have been proposed in [9], [10], [11] and [14]. The majority of the work related to data replication on the Internet employs the coarse-grained (site based) replication. As the Internet grows and the limitations of caching become more obvious, the importance of fine-grained replication, *i.e.*, duplicating highly popular data objects, is likely to increase [11].

In this paper, the replica schemas are established in a static fashion. The aim is to identify a replica schema that effectively minimizes the object transfer cost. We propose a novel technique based on the extension of Vickrey auction [17] called the  $N+1^{\text{st}}$  price auction mechanism (NPAM), where the players

compete for memory space at sites so that replicas can be placed. This approach is compared against three well-known techniques from the literature: branch and bound [11], greedy [15], and genetic algorithms [13]. Experimental results reveal that this simple and intuitive approach outperforms the three techniques in both execution time and solution quality.

The remainder of this paper is organized as follows. Section 2 provides the related work. Section 3 formulates the DRP. Section 4 concentrates on modeling the resource allocation mechanism for the DRP. The experimental results and concluding remarks are provided in Sections 5 and 6, respectively.

## 2 Related Work

The data replication problem (see Section 3 for a formal description) is an extension of the classical file allocation problem (FAP). Chu [4] studied the file allocation problem with respect to multiple files in a multiprocessor system. Casey [4] extended this work by distinguishing between updates and read file requests. Eswaran [7] proved that Casey's formulation was NP complete. In [14] Mahmoud *et al.* provide an iterative approach that achieves good solution quality when solving the FAP for infinite server capacities. A complete although old survey on the FAP can be found in [6]. Apers in [1] considered the data allocation problem (DAP) in distributed databases where the query execution strategy influences allocation decisions. In [11] the authors proposed several algorithms to solve the data allocation problem in distributed multimedia databases (without replication), also called as video allocation problem (VAP).

Most of the research papers outlined in [6] aim at formalizing the problem as an optimization one, sometimes using multiple objective functions. Network traffic, server throughput and response time exhibited by users are considered for optimization. Although a lot of effort was devoted in providing comprehensive models, little attention has been paid to good heuristics for solving this complex problem. Furthermore access patterns are assumed to remain static and solutions in the dynamic case are obtained by re-

executing a time consuming mathematical programming technique.

Some on-going work is related to dynamic replication of objects in distributed systems when the read-write patterns are not known *a priori*. Awerbuch's *et al.* work in [2] is significant from a theoretical point of view, but the adopted strategy for commuting updates (object replicas are first deleted), can prove difficult to implement in a real-life environment. In [16] Wolfson *et al.* proposed an algorithm that leads to optimal single file replication in the case of a tree network. The performance of the scheme for general network topologies is not clear though. Dynamic replication protocols were also considered under the Internet environment. In [15], Rabinovich *et al.* proposed a protocol for dynamically replicating the contents of an ISP (Internet Service Provider) in order to improve client-server proximity without overloading any of the servers. However updates were not considered.

## 3 Problem Formulation

Consider a distributed system comprising  $M$  sites, with each site having its own processing power, memory (primary storage) and media (secondary storage). Let  $S_i$  and  $s_i$  be the name and the total storage capacity (in simple data units e.g. blocks), respectively, of site  $i$  where  $1 \leq i \leq M$ . The  $M$  sites of the system are connected by a communication network. A link between two sites  $S_i$  and  $S_j$  (if it exists) has a positive integer  $c(i,j)$  associated with it, giving the communication cost for transferring a data unit between sites  $S_i$  and  $S_j$ . If the two sites are not directly connected by a communication link then the above cost is given by the sum of the costs of all the links in a chosen path from site  $S_i$  to the site  $S_j$ . Without the loss of generality we assume that  $c(i,j) = c(j,i)$ . This is a common assumption (e.g. see [10], [11], and [14]). Let there be  $N$  objects, each identifiable by a unique name  $O_k$  and size in simple data unites  $o_k$  where  $1 \leq k \leq N$ . Let  $r_k^i$  and  $w_k^i$  be the total number of reads and writes, respectively, initiated from  $S_i$  for  $O_k$ .

Our replication policy assumes the existence of one primary copy for each object in the network. Let  $P_k$ , be the site which holds the primary copy of  $O_k$ , *i.e.*, the only copy in the network that

cannot be de-allocated, hence referred to as primary site of the  $k$ -th object. Each primary site  $P_k$ , contains information about the whole replication scheme  $R_k$  of  $O_k$ . This can be done by maintaining a list of the sites where the  $k$ -th object is replicated at, called from now on the *replicators* of  $O_k$ . Moreover, every site  $S_i$  stores a two-field record for each object. The first field is its primary site  $P_k$  and the second the nearest neighborhood site  $NN_k^i$  of site  $S_i$  which holds a replica of object  $k$ . In other words,  $NN_k^i$  is the site for which the reads from  $S_i$  for  $O_k$ , if served there, would incur the minimum possible communication cost. It is possible that  $NN_k^i = S_i$ , if  $S_i$  is a *replicator* or the primary site of  $O_k$ . Another possibility is that  $NN_k^i = P_k$ , if the primary site is the closest one holding a replica of  $O_k$ . When a site  $S_i$  reads an object, it does so by addressing the request to the corresponding  $NN_k^i$ . For the updates we assume that every site can update every object. Updates of an object  $O_k$  are performed by sending the updated version to its primary site  $P_k$ , which afterwards broadcasts it to every site in its replication scheme  $R_k$ .

For the DRP under consideration, we are interested in minimizing the total Replication Cost (RC) (or the total network transfer cost) due to object movement, since the communication cost of control messages has minor impact to the overall performance of the system. There are two components affecting RC. The first component of RC is due to the read requests. Let  $R_k^i$  denote the total RC, due to  $S_i$ 's reading requests for object  $O_k$ , addressed to the nearest site  $NN_k^i$ . This cost is given by the following equation:

$$R_k^i = r_k^i o_k c(i, NN_k^i), \quad (1)$$

where  $NN_k^i = \{Site\ j \mid j \in R_k \wedge \min c(i, j)\}$ . The second component of RC is the cost arising due to the writes. Let  $W_k^i$  be the total RC, due to  $S_i$ 's writing requests for object  $O_k$ , addressed to the primary site  $P_k$ . This cost is given by the following equation:

$$W_k^i = w_k^i o_k \left( c(i, P_k) + \sum_{\forall j \in R_k, j \neq i} c(NN_k^i, j) \right). \quad (2)$$

Here, we made the indirect assumption that in order to perform a write we need to ship the whole updated version of the object. This of course is not always the case, as we can move only the updated parts of it (modeling such

policies can also be done using our framework). The cumulative RC, denoted as  $C_{overall}$ , due to reads and writes is given by:

$$C_{overall} = \sum_{i=1}^M \sum_{k=1}^N (R_k^i + W_k^i). \quad (3)$$

Let  $X_{ik} = 1$  if  $S_i$  holds a replica of object  $O_k$ , and 0 otherwise.  $X_{ik}$ s define an  $M \times N$  replication matrix, named  $X$ , with boolean elements. Equation 3 is now refined to:

$$X = \sum_{i=1}^M \sum_{k=1}^N \left[ \begin{array}{l} (1 - X_{ik}) \left[ r_k^i o_k \min \{c(i, j) \mid X_{jk} = 1\} \right. \\ \left. + w_k^i o_k c(i, P_k) \right] + X_{ik} \left( \sum_{s=1}^M w_s^i \right) o_k c(i, P_k) \end{array} \right]. \quad (4)$$

Sites which are not the *replicators* of object  $O_k$  create RC equal to the communication cost of their reads from the nearest *replicator*, plus that of sending their writes to the primary site of  $O_k$ . Sites belonging to the replication scheme of  $O_k$ , are associated with the cost of sending/receiving all the updated versions of it. Using the above formulation, the DRP can be defined as:

“Find the assignment of 0,1 values in the  $X$  matrix that minimizes  $C_{overall}$ , subject to the storage capacity constraint:

$$\sum_{k=1}^N X_{ik} o_k \leq s_i \quad \forall (1 \leq i \leq M),$$

and subject to the primary copies policy:

$$X_{P_k k} = 1 \quad \forall (1 \leq k \leq N).”$$

The minimization of  $C_{overall}$  has the following two impacts on the distributed system under consideration. First, it ensures that the object replication is done in such a way that it minimizes the maximum distance between the replicas and their respective primary objects. Second, it ensures that the maximum distance between an object  $k$  and the user(s) accessing that object is also minimized. Thus, the solution aims for reducing the overall RC of the system. In the generalized case, the DRP is NP-complete [13].

## 4 The Mechanism (NPAM)

We term the proposed resource allocation mechanism as NPAM an acronym for N+1<sup>st</sup> Price Auction Mechanism. In the auction setup each primary copy of an object  $k$  is a player. A player  $k$  can perform the necessary computations on its strategy set by using the site (where it resides)  $P_k$ 's processor. At each given instance a (sub)-auction takes place at a particular site  $i$  chosen in

---

```

N+1st Price Auction Mechanism
Initialize:
01  $LS, L^i$ .
02 WHILE  $LS \neq \text{NULL}$  DO
03   SELECT  $S^j \in LS$  /*Round-robin fashion */
04   FOR each  $k \in O$  DO
05      $B_k = \text{compute}(B_k^j)$ ; /*compute the benefit*/
06     Report  $B_k$  to  $S_j$  which stores in array  $B$ ;
07   END FOR
08   WHILE  $b_i \geq 0$ 
09      $B_k = \text{argmax}_k(B)$ ; /*Choose the best offer*/
10     Extract the info from  $B_k$  such as  $O_k$  and  $o_k$ ;
11      $b_i = b_i - o_k$ ; /*Calculate space and termination condition*/
12     Payment =  $B_k$ ; /* Maintain N+1st price */
13     IF  $b_i < 0$  THEN EXIT WHILE ELSE
14        $L^i = L^i - O_k$ ; /*Update the list*/
15       Update  $NN_{OMAX}$  /*Update the nearest neighbor list*/
16       IF  $L^i = \text{NULL}$  THEN SEND info to  $M$  to update  $LS = LS - S^i$ ;
17       Replicate  $O_k$ ;
18     END WHILE
19      $S_j$  asks all successful bidders to pay  $B_k$ 
20 END WHILE

```

---

FIGURE I  
PSEUDO-CODE FOR N+1<sup>ST</sup> PRICE AUCTION MECHANISM.

a round robin fashion from the set of  $M$  sites. These auctions are performed continuously throughout the system's life, making it a self evolving and self repairing system. However, for simulation purposes ("cold" network [11]) we discrete the continuum solely for the reason to observe the solution quality.

Each player  $k$  competes through bidding for memory at a site  $i$ . Many would argue that memory constraints are no longer important due to the reduced costs of memory chips. However, replicated objects (just as cached objects) reside in the memory (primary storage) and not in the media (secondary storage) [13]. Thus, there will always be a need to give priority to objects that have higher access (read and write) demands. Moreover, memory space regardless of being primary or secondary is limited.

Each player  $k$ 's strategy is to place a replica at a site  $i$ , so that it maximizes its (the object's) benefit function. The benefit function gives more weight to the objects that incur reduced RC in the system:

$$B_k^i = R_k^i - \left( \sum_{x=1}^M w_k^x o_k^x c(i, P_k) - W_k^i \right). \quad (5)$$

The above value represents the expected benefit (in RC terms), if  $O_k$  is replicated at  $S_i$ . This benefit is computed using the difference between the read and update cost. Negative values of  $B_k^i$  mean that replicating  $O_k$  is inefficient from the "local view" of  $S_i$  (although it

might reduce the global RC due to bringing the object closer to other servers). The pseudo-code for the N+1<sup>st</sup> price auction is given in Figure 1.

We maintain a list  $L^i$  at each server. The list contains all the objects that can be replicated at  $S_i$  (i.e., the remaining storage capacity  $b^i$  is sufficient and the benefit value is positive). We also maintain a list  $LS$  containing all servers that can replicate an object. In other words,  $S_i \in LS$  if and only if  $L^i \neq \text{NULL}$ . The auction mechanism performs in steps. In each step a server  $S_i$  is chosen from  $LS$  in a round-robin fashion. Each player  $k \in O$  calculates the benefit function of object. The set  $O$  represents the collection of players that are legible for participation. A player  $k$  is legible if and only if the benefit function value obtained for site  $S_i$  is the maximum of among all the other benefit function values for sites other than  $i$ , i.e.,  $S_i \geq S_{-i}$ . This is done in order to suppress mediocre bids, which, in turn improves computational complexity. It is to be noted that in each step  $L^i$  together with the corresponding nearest server value  $NN_k^i$ , are updated accordingly.

*Theorem 1:* NPAM takes  $O(MN^2)$  time.

*Proof:* The worst case execution time of the algorithm is when each server has sufficient capacity to store all objects and the update ratios are low enough so that no object incurs negative benefit value. In that case, the while-loop (02) performs  $M$  iterations. The time complexity for each iteration is governed by the for-loop in (04) and the while loop in (08) ( $O(N^2)$  in total). Hence, we conclude that the worst case running time of the algorithm is  $O(MN^2)$ . ■

## 5 Experimental Results

We performed experiments on a 440MHz Ultra 10 machine with 512MB memory. The experimental evaluations were targeted to benchmark the placement policies. The solution quality in all cases, was measured according to the RC percentage that was saved under the replication scheme found by the algorithms, compared to the initial one, i.e., when only primary copies exist.

To evaluate our proposed technique on realistic traffic patterns, we used the access logs collected at the Soccer World Cup 1998 website [2]. Each

TABLE II  
RUNNING TIME IN SECONDS

Problem Size	Greedy	GRA	Aε-Star	NPAM
M= 500, N= 1350	<b>81.69</b>	117.60	110.46	90.09
M= 500, N= 1400	98.28	127.89	127.89	<b>95.34</b>
M= 500, N= 1450	122.43	139.02	139.02	<b>98.91</b>
M= 500, N= 1500	134.61	148.47	155.40	<b>104.37</b>
M= 500, N= 1550	146.58	168.84	169.47	<b>105.63</b>
M= 500, N= 2000	152.25	177.66	189.21	<b>108.57</b>

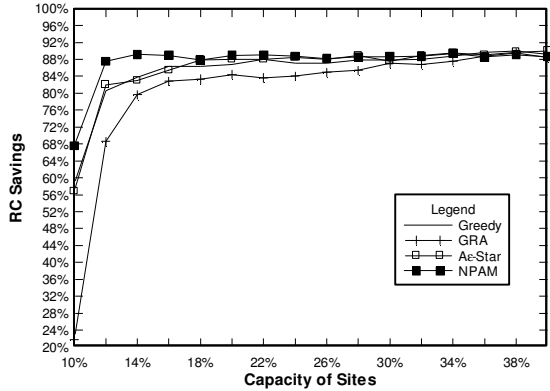


FIGURE II  
RC SAVINGS VS. SYSTEM CAPACITY (N=2000, M=500, U=5%).

experimental setup was evaluated thirteen times, *i.e.*, Friday (24 hours) logs from May 1, 1998 to July 24, 1998. To process the logs, we wrote a script that returned: only those objects which were present in all the logs (2000 in our case), the total number of requests from a particular client for an object, the average and the variance of the object size. From this log we choose the top five hundred clients (maximum experimental setup), which were randomly mapped to one of the nodes of the topologies. Note that this mapping is not 1-1, rather 1- $M$ . This gave us enough skewed workload to mimic real world scenarios. It is also worthwhile to mention that the total amount of requests entertained for each problem instance was in the range of 1-2 million. The primary replicas' original site was mimicked by choosing random locations. The capacities of the sites  $C\%$  were generated randomly with range from *Total Primary Object Sizes*/2 to  $1.5 \times$  *Total Primary Object Sizes*. The variance in the object size collected from the access logs helped to instill enough diversity to benchmark object updates. The updates were randomly pushed onto different sites, and the total system update load was measured in terms of the percentage update requests  $U\%$  compared that to the initial network




FIGURE III  
RC SAVINGS VS. READS (N=2000, M=500, C=45%).




FIGURE IV  
RC SAVINGS VS. UPDATES (N=2000, M=500, C=60%).

with no updates.

For comparison, we selected three various types of replica placement techniques. To provide a fair comparison, the assumptions and system parameters were kept the same in all the approaches. We chose 1) from [11] the efficient branch-and-bound based technique (Aε-Star), 2) from [13] the genetic algorithm based technique (GRA) which showed excellent adaptability against skewed workload, 3) and from [15] the famous greedy approach (Greedy). Due to space limitations, we briefly describe the comparative approaches. Details for a specific technique can be obtained from the referred papers.

Table II (best times shown in bold) shows the algorithm execution times. The number of sites was kept constant at 500, and the number of objects was varied from 1350 to 2000. With maximum load (2000 objects and 500 sites), the

TABLE III  
AVERAGE RC SAVINGS IN PERCENTAGE

Problem Size	Greedy	GRA	A $\epsilon$ -Star	NPAM
$N=150, M=20 [C=20\%, U=25\%]$	70.46	69.74	74.62	<b>75.70</b>
$N=200, M=50 [C=20\%, U=20\%]$	73.94	70.18	77.42	<b>78.43</b>
$N=300, M=50 [C=25\%, U=5\%]$	70.01	64.29	70.33	<b>82.25</b>
$N=300, M=60 [C=35\%, U=5\%]$	71.66	65.94	72.01	<b>74.43</b>
$N=400, M=100 [C=25\%, U=25\%]$	67.40	62.07	71.26	<b>73.89</b>
$N=500, M=100 [C=30\%, U=35\%]$	66.15	61.62	71.50	<b>75.45</b>
$N=800, M=200 [C=25\%, U=15\%]$	67.46	65.91	70.15	<b>73.68</b>
$N=1000, M=300 [C=25\%, U=35\%]$	69.10	64.08	70.01	<b>72.45</b>
$N=1500, M=400 [C=35\%, U=50\%]$	70.59	63.49	70.51	<b>74.01</b>
$N=2000, M=500 [C=10\%, U=60\%]$	67.03	63.37	72.16	<b>73.15</b>

proposed technique NPAM saved approximately 50 seconds of termination time then the second fastest algorithm (Greedy).

Superiority of execution time comes at the cost of loss in solution quality. However, NPAM showed high solution quality. First, we observe the effects of system capacity increase. An increase in the storage capacity means that a large number of objects can be replicated. Replicating an object that is already extensively replicated, is unlikely to result in significant traffic savings as only a small portion of the servers will be affected overall. Moreover, since objects are not equally read intensive, increase in the storage capacity would have a great impact at the beginning (initial increase in capacity), but has little effect after a certain point, where the most beneficial ones are already replicated. This is observable in Figure II, which shows the performance of the algorithms. Greedy and NPAM showed an immediate initial increase (the point after which further replicating objects is inefficient) in its RC savings, but afterward showed a near constant performance. GRA although performed the worst, but observably gained the most RC savings (35%) followed by Greedy with 29%. Further experiments with various update ratios (5%, 10%, and 20%) showed similar plot trends. It is also noteworthy (plots not shown in this paper due to space restrictions) that the increase in capacity from 10% to 17%, resulted in 4 times (on average) more replicas for all the algorithms.

Next, we observe the effects of increase in the read and update (write) frequencies. Since these two parameters are complementary to each other, we describe them together. In both the setups the number of sites and objects were kept constant. Increase in the number of reads in the system would mean that there is a need to replicate as

many object as possible (closer to the users). However, the increase in the number of updates in the system requires the replicas be placed as close as to the primary site as possible (to reduce the update broadcast). This phenomenon is also interrelated with the system capacity, as the update ratio sets an upper bound on the possible traffic reduction through replication. Thus, if we consider a system with unlimited capacity, the “replicate everywhere anything” policy is strictly inadequate. The read and update parameters indeed help in drawing a line between good and marginal algorithms. The plots in Figures III and IV show the results of read and update frequencies, respectively. A clear classification can be made between the algorithms. A $\epsilon$ -Star, Greedy and NPAM incorporate the increase in the number of reads by replicating more objects and thus savings increase up to 89%. GRA gained the least of the RC savings of up to 67%. To understand why there is such a gap in the performance between the algorithms, we recall from [13] that GRA specifically depends on the initial population of the candidate solution. Moreover, GRA maintains a localized network perception. Increase in updates result in objects having decreased local significance (unless the vicinity is in close proximity to the primary location). On the other hand, A $\epsilon$ -Star, Greedy and NPAM never tend to deviate from their global view of the problem domain.

In summary, Table III shows the quality of the solution in terms of RC percentage for 10 problem instances (randomly chosen), each being a combination of various numbers of sites and objects, with varying storage capacity and update ratio. For each row, the best result is indicated in bold. The proposed NPAM steals the show in the context of solution quality, but A $\epsilon$ -Star and Greedy do indeed give a good competition, with savings within a range of 7%-10% of NPAM.

## 6 Conclusions

Manual mirroring of data objects is a tedious and time consuming operation. This paper proposed a game theoretical  $N+1^{\text{st}}$  price auction mechanism (NPAM) for fine-grained data replication in large-scale distributed computing systems such as the Internet. NPAM is a protocol for automatic

replication and migration of objects in response to demand changes. NPAM aims to place objects in the proximity of a majority of requests while ensuring that no hosts become overloaded.

NPAM allows agents to compete for the scarce memory space at sites so that they can acquire the rights to place replicas. To cater for the possibility of cartel type behavior of the agents, NPAM uses N+1<sup>st</sup> price protocol. This leaves the agents with no option, then to report truthful valuations of the objects that they represent.

NPAM was compared against some well-known techniques, such as: greedy, branch and bound and genetic algorithms. To provide a fair comparison, the assumptions and system parameters were kept the same in all the approaches. The experimental setup was designed to mimic a large-scale distributed computing system (the Internet), by using several Internet topology generators and World Cup Soccer 1998 web server access logs. The experimental results revealed that NPAM outperformed the three widely cited and powerful techniques in both the execution time and solution quality. In summary, NPAM exhibited 7%-10% better solution quality and 10%-30% savings in the algorithm termination timings.

## References

[1] P. Apers. "Data Allocation in Distributed Database Systems". *ACM Trans. Database Systems*. 13(3): 263-304, 1988.

[2] M. Arlitt and T. Jin. "Workload characterization of the 1998 World Cup Web Site". Tech. report. HP Lab. Palo Alto. HPL-1999-35(R.1). 1999.

[3] B. Awerbuch, Y. Bartal and A. Fiat. "Competitive Distributed File allocation". *Proc. 25th ACM STOC*. 1993, pp. 164-173.

[4] K. Calvert, M. Doar, E. Zegura. "Modeling Internet topology". *IEEE Communications*, 35(6): 160-163, 1997.

[5] R. Casey. "Allocation of Copies of a File in an Information Network". *Proc. Spring Joint Computer Conf.* 1972, pp. 617-625.

[6] W. Chu. "Optimal File Allocation in a Multiple Computer System". *IEEE Trans. on*

*Computers*. C-18(10): 885-889, 1969.

[7] L. Dowdy and D. Foster. "Comparative Models of the File Assignment problem". *ACM Computing Surveys*. 14(2): 287-313, 1982.

[8] K. Eswaran. "Placement of Records in a File and File Allocation in a Computer Network". *Information Processing Letters*. 1: 304-307, 1974.

[9] S. Floyd and V. Paxson. "Difficulties in simulating the internet". *IEEE/ACM Transactions on Networking*. 9(4): 253-285, 2001.

[10] J. Kangasharju, J. Roberts and K. Ross. "Object Replication Strategies in Content Distribution Networks". *Proc. of WCCD*. 2001, pp. 455-466.

[11] S. Khan and I. Ahmad. "Heuristic-based Replication Schemas for Fast Information Retrieval over the Internet". *Proc. of 17th International Conference on Parallel and Distributed Computing*. 2004, pp. 278-283.

[12] Y. Kwok, K. Karlapalem, I. Ahmad and N. Pun. "Design and Evaluation of Data Allocation Algorithms for Distributed Database Systems". *IEEE Journal on Selected areas in Communication*. 14(7): 1332-1348, 1996.

[13] T. Loukopoulos and I. Ahmad. "Static and Adaptive Distributed Data Replication using Genetic Algorithms". *Journal of Parallel and Distributed Computing*, 64(11): 1270-1285, 2004.

[14] S. Mahmoud and J. Riordon. "Optimal Allocation of Resources in Distributed Information Networks," *ACM Trans. on Database Systems*. 1(1) 66-78, 1976.

[15] L. Qiu, V. Padmanabhan and G. Voelker. "On the Placement of Web Server Replicas". *Proc. of the IEEE INFOCOM*, 2001.

[16] M. Rabinovich. "Issues in Web Content Replication," *Data Engineering Bulletin*. 21(4): 21-29, 1998.

[17] W. Vickrey, "Counter-speculations, Auctions, and Competitive Sealed-bid Tenders". *Journal of Finance*. (16):8-37, 1961.

[18] O. Wolfson, S. Jajodia and Y. Hang. "An Adaptive Data Replication Algorithm". *ACM Trans. on Database Systems*. 22(4):255-314, 1997.

[19] G. Zipf. "Human Behavior and the Principle of Least-Effort". Addison-Wesley. 1949.